

SYSTEMS AND METHODS FOR
NEGOTIATED RESOURCE UTILIZATION

This application claims the benefit of U.S. provisional application number 60/247184
5 incorporated herein by reference in its entirety.

Field of The Invention

The field of the invention is resource load balancing.

Background of The Invention

With the increasing number of users on the Internet comes the increasing probability
10 that a single server will become inundated with requests for a resource. The result is often
slow response times, large queues, and even system failure, despite the probable availability
of alternate servers that could aid the primary server by off loading some of its tasks. This
creates a problem of both resource over-utilization and resource under-utilization, both of
which result in negative financial impact.

15 To alleviate these issues, an art known as 'load balancing' has been developed. Local
load balancing systems are known, mostly based on hardware switches that route resource
requests to one of a plurality of mirrored (i.e. containing identical content/resources) servers,
giving the external appearance of a single, high-powered server. Global server load
balancing systems have also recently emerged that attempt to distribute the traffic directed to
20 any given server to a plurality of servers, which are geographically distributed. See, for
example, the approaches of Akamai (www.akamai.com) and Digital Island
(www.digisite.net) for the distribution of web traffic.

As an alternative motivating example, consider a large data processing operation that
has computing facilities in numerous locations throughout the world. If one location is
25 overwhelmed, a systems administrator (whether human or machine) may choose to route
additional operations to other locations. Similarly, it may be advantageous to employ
several computers in a parallel processing mode to accomplish a single task. Even though

several computers may be utilizing resources on other computers, there is never any negotiation among the several computers as to what tasks each one will take, and under what terms. If a computer is instructed to perform a given task, it takes on that task. The computer cannot refuse the work because the subject matter is undesirable for some reason,
5 or the compensation is not sufficient.

All previously known load balancing systems and methods operate on a master-slave configuration. That scenario may work well for systems that are all owned or operated by the same company, or at least under the umbrella of a single company or other "master", but doesn't work well at all for computers owned by independent operators. Yet there are many
10 millions of independent computers throughout the world that are incredibly under-utilized, usually with under 20% utilization, and often with utilization of less than 5% or even 1%. While such users may at times make their computers available for such projects as the Search for Extraterrestrial Intelligence (SETI - see www.seti.org for details of their distributed processing utilizing screen savers), those computers, and their associated
15 processing, storage, and other resources, will likely never be made available to others unless there are some sort of systems and methods in place by which the use of those resources can be automatically negotiated.

Consequently, there is a need in the art for systems where one computer can automatically recruit resources controlled by another, independent, computer through a
20 negotiation between the computers.

The present invention provides systems and methods in which a resource is recruited by a negotiation between two independent computers. The computers are preferably at a significant distance from one another, and likely not part of the same local area network (LAN). Interaction of more than two computers is also contemplated, such that a single
25 computer may negotiate resource usage with more than one other computer, and computers could act as consolidators/brokers to combine and offer for usage the combination of resources from multiple other computers. Contemplated resources are very wide-ranging, and include one or more processors, printers, storage devices, and communication channels.

The negotiation may be based on any condition or combination of conditions,

5

10

Summary of the Invention

by a negotiation between two independent computers.

15

20

25

The negotiation may be facilitated by one of the computers listing either a need or an available resource on a directory. To that end the directory may advantageously include a plurality of resource descriptions, at least one condition for use of each of the plurality of resource descriptions, and at least one item of ratings information for each of the plurality of resource descriptions.

Various objects, features, aspects, and advantages of the present invention will become more apparent from the following detailed description that describes a preferred embodiment of the invention, along with the accompanying drawings in which like numerals represent like components.

10 **Brief Description of The Drawing**

Figure 1 is a schematic of a system by which a requesting computer negotiates use of a resource controlled by a providing computer.

Figure 2 is a table representation of data in a resource directory.

Detailed Description

15 In Figure 1, a system 5 generally comprises a requesting computer 10 communicating with a providing computer 30 to obtain access to a resource 35. Optionally (a) a directory hosting computer 30 may host a resource directory 22, (b) the providing computer 30 may provide access to other resources 36, 37, 38, and (c) an additional providing computer 40 may provide access to additional resources 45, 46, 47.

20 Communication paths 11, 12, 13, 31, 32 provide one or two way connectivity among the various elements as shown.

Nature of the Computers

Requesting computer 10, and indeed all of the various computers described herein, may be any sort of computing device or system. Thus, computer 10 may be a
25 supercomputer or other massively parallel device, a mainframe, minicomputer, desktop, laptop, personal digital assistant (PDA), or other microcomputer, as well as specialty devices

5

10

15

20

25

Recruitment Message

The requesting computer 10 may advantageously monitor itself with respect to efficiency, backlog, or other operations parameter or parameters. Such monitoring may be historical, concurrent, or even proactive, i.e., planning for future bottlenecks depending upon anticipated workload requirements. When the requesting computer 10 detects a condition of over utilization, inefficiency, or some other undesirable situation involving at least one of the parameters, it generates a recruitment message (not shown) that travels along communication path 12 to the resource computer 30.

In one contemplated embodiment, the requesting computer 10 multicasts the recruitment message for assistance to a plurality of computers until it is able to get assistance. In an alternative embodiment, the requesting computer 10 posts its requirement on some sort of bulletin board or other directory, as for example the directory 22 operated by the directory hosting computer 20.

The recruitment message preferably includes information sufficient for a potential providing computer to determine if a match is realistic. For example, if the requesting computer 10 is willing to pay only a much smaller amount for unit of resource than the offering computer is willing to accept, then there may be no need for negotiations. Similarly, if a potential providing computer does not have the resource needed, or can gain access to such resource at a reasonably time and cost, then there may be no need for negotiations. Sufficient information presently contemplated in at least some embodiments includes a priority for the task, the type of resource needed, the type of task, and a deadline for completion of the task.

It should be appreciated that where the recruitment message includes the type of task, a further benefit may be derived by delegating certain tasks to relatively slow processors. Some tasks do not require a high speed CPU to perform them, and therefore they may be delegated to one or more relatively slow processors without any significantly undesirable effects. This may reduce obsolescence of an enormous number of machines.

It should also be appreciated that the providing computer 20 may advantageously

monitors itself in a similar manner to the requesting computer 10. If it detects a condition of under utilization for the resource, it may send out a message identifying the resources available, and the pricing or other conditions under which that resource may be utilized.

Communication Paths

5 The path 12 may be determined by any combination of direct addressing, active packet routing, narrow cast, and broadcast. The resource computer 30 may be selected by any suitable method, including an historical experience of the requesting computer 10 in utilizing a resource of the providing computer 30. Alternatively or additionally, the requesting computer 10 may take advantage of the historical experience of some other
10 computer (not shown) in utilizing a resource of the providing computer 30. One contemplated embodiment of that scenario may involve posting of information relating to the historical experience on a directory such as directory 22. In yet another alternative embodiment, the resource computer 30 may be selected by the resource computer 30 sending a response (not shown) to the recruitment message.

15 The various communication paths 11, 12, 13, 31, 32 may be embodied in any suitable manner. Very likely one or more of such paths can take place on a public, packet switched network such as the Internet, and in such instances the paths would likely not be dedicated. On the other hand, all suitable embodiments are contemplated. It is contemplated, for example, that one or more of the paths 11, 12, 13, 31, 32 may be carried
20 on a LAN or Wide Area Network other than the Internet. From another perspective, it is contemplated that one or more of the paths 11, 12, 13, 31, 32 may take place using hard wired telephone lines, fiber optics or other cabling, or some sort of local or global wireless transmission.

 It is important to note that communication paths 11, 12, 13, 31, 32 need not be
25 utilized for transfer of programs, data, or results. The various computers 10, 30, 40 may well utilize other paths (not shown) instead of, or in addition to the paths defined in Figure 1.

Negotiation

The requesting computer 10 and the providing computer 30 negotiate at least one condition under which the providing computer 30 provides a resource 35 to the requesting computer. As used herein, the term "negotiate" means as a "bargained for" exchange.

5 Consequently, no one party to the negotiation completely controls the terms. Also as used herein, the term "providing the resource" does not necessarily mean that the entire resource is turned over to the requesting computer 10, and indeed it is more likely that only a portion of the resource 35 will be made available to the requesting computer 10. Thus, if the resource 35 is a mass storage device such as a hard disk or collection of hard disks such as
10 may be found in an EMC™ mass storage system, it is very likely that only a portion of the hard disk or mass storage system will be utilized by the requesting computer 10.

Although the negotiation can occur in any suitable manner, it is preferred that the negotiation occur in a completely automated manner, i.e. without direct human intervention.

Presumably this means an exchange of messages between the requesting computer 10 and
15 the providing computer 30. It is, however, also contemplated that negotiation can be conducted by, or at least assisted by, some sort of human or machine agent (not shown).

The requesting computer 10 and the providing computer 30 preferably negotiate based upon established negotiating algorithms, the conditions negotiated include the amount of compensation, the subject matter of the task, and the ability to perform the work.

20 Often, the result of the negotiation is formation of an "enclave" to perform a defined task that is to be completed by a deadline. Alternatively or additionally, the negotiation may result in performance of relatively undefined tasks, which may continue for days, weeks, or even longer periods of time. It is appreciated that formation and maintenance of the enclave may well involve many issues besides those discussed herein, including security
25 considerations, trust relationships, adding and deleting members of the enclave, work distribution, and so forth.

Directory

The system 5 optionally includes a directory hosting computer 20 that hosts a

directory 22. In this case the directory 22 is designed to maintain a list of available resources, including many conditions. A more detailed explanation is set forth below with respect to Figure 2. Directory hosting computer 20 could additionally or alternatively include a list of requests.

5 One aspect of the directory 22 is that it may include peer ratings resulting from prior usages of the resources. Peer ratings involves a subsystem where computers maintain historical statistics about other computers they have received assistance from. Contemplated statistics comprise an overall satisfaction value, a type of task, a reliability ratings, and a promptness ratings. One purpose of such a peer ratings system would be to provide a base
10 line that the requesting computer 10 can use in its negotiation with the provider computer 30. If the ratings of the providing computer 20 with respect to resource 35 is relatively poor, the requesting computer 10 may well negotiate more favorable terms. Thus, the rate (in money or otherwise) that the requesting computer 10 may pay the providing computer may at least partially be based upon the peer ratings provided by the directory hosting computer
15 30.

Methods

Figure 1 can be understood in method terms. In a broad interpretation, Figure 1 depicts a method of recruiting a resource 35, comprising: establishing a communications link 12 between a requesting computer 10 and a providing computer 30 that is not controlled
20 by the requesting computer 10; and the requesting computer 10 and the providing computer 30 negotiating a condition under which the providing computer 30 provides the resource 35 to the requesting computer 10.

The requesting computer is preferably located at a distance of at least 1 kilometer from the providing computer. If the requesting computer 10 and providing computer 30 are
25 members of local area networks (not shown), they are preferably not members of the same LAN.

The step of negotiating preferably comprises negotiating a compensation rate for the use of the resource, which may be monetary, quid-pro-quo, or may involve some other

5

10

15

20

25

The order and indeed to exact content of the various columns is quite variable. In this instance column 101 contains the type of available resource, column 102 contains the time of day at which the resources are available, column 103 contains a restriction on subject matter, column 104 contains a priority that a task will be allotted, column 105 contains the compensation demanded, column 106 contains the likelihood of downtime, column 107 contains an aggregate satisfaction value, columns 108 - 111 (not shown) contain recency values, reliability values, type of work values, and IP address of the providing computers, respectively, and column 112 contains promptness values. Other contemplated tables could readily have a greater or lesser number of columns.

Thus, specific embodiments and applications of methods and apparatus of the present invention have been disclosed. It should be apparent, however, to those skilled in the art that many more modifications besides those described are possible without departing from the inventive concepts herein. The inventive subject matter, therefore, is not to be restricted except in the spirit of the appended claims.